

Published in final edited form as:

Nat Methods. 2020 August 01; 17(8): 799–806. doi:10.1038/s41592-020-0885-x.

Investigating higher order interactions in single cell data with scHOT

Shila Ghazanfar¹, Yingxin Lin², Xianbin Su³, David Ming Lin⁴, Ellis Patrick^{2,5}, Ze Guang Han³, John C. Marioni^{#1,6}, Jean Yee Hwa Yang^{#2,7,*}

¹Cancer Research UK Cambridge Institute, University of Cambridge, Cambridge, United Kingdom

²School of Mathematics and Statistics, The University of Sydney, Sydney, New South Wales, 2006, Australia

³Key Laboratory of Systems Biomedicine (Ministry of Education) and Collaborative Innovation Center of Systems Biomedicine, Shanghai Center for Systems Biomedicine, Shanghai Jiao Tong University, Shanghai 200240, China

⁴Department of Biomedical Sciences, Cornell University, Ithaca, United States

⁵Westmead Institute for Medical Research, Sydney, New South Wales, 2145, Australia

⁶European Molecular Biology Laboratory, European Bioinformatics Institute, (EMBL-EBI), Wellcome Genome Campus, Cambridge, United Kingdom

⁷Wellcome Sanger Institute, Wellcome Genome Campus, Cambridge, United Kingdom

⁷Charles Perkins Centre, The University of Sydney, Sydney, New South Wales, 2006, Australia

These authors contributed equally to this work.

Abstract

Single-cell genomics has transformed our ability to examine cell fate choice. Examining cells along a computationally ordered “pseudotime” offers the potential to unpick subtle changes in variability and covariation among key genes. We describe a novel approach, scHOT – single cell Higher Order Testing - which provides a flexible and statistically robust framework for identifying changes in higher order interactions among genes. scHOT can be applied for cells along a continuous trajectory or across space and accommodates any higher order measurement including variability or correlation. We demonstrate the utility of scHOT by studying coordinated changes in higher order interactions during embryonic development of the mouse liver. Additionally, scHOT identifies subtle changes in gene-gene correlations across space using spatially-resolved

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use: http://www.nature.com/authors/editorial_policies/license.html#terms

*To whom correspondence should be addressed.; marioni@ebi.ac.uk; jean.yang@sydney.edu.au.

Author contributions

S.G. conceived the study with input from E.P., J.Y.H.Y. and J.C.M.. S.G. developed the method and software and performed data analysis with input from Y.L., S.G., D.M.L. and X.S. interpreted the results with input from Z.G.H. and J.C.M.. S.G., J.C.M. and J.Y.H.Y. wrote the manuscript. All authors read and approved the final version of the manuscript.

Ethics declaration

The authors declare no competing interests.

transcriptomics data from the mouse olfactory bulb. scHOT meaningfully adds to first order differential expression testing and provides a framework for interrogating higher order interactions using single cell data.

Introduction

Understanding the mechanisms that underpin cell fate choice is a key challenge in developmental biology. It requires disentangling the complex interplay between cell autonomous factors, such as gene expression, and non-autonomous factors such as the signaling environment. In the former context, recent technological advances have enabled the rapid and high-throughput measurement of mRNA expression levels in individual cells. Such single-cell RNA-sequencing (scRNA-seq) datasets have facilitated the generation of atlases of cell types during development in human, mouse, zebrafish and the frog^{1–5}. Using such data, cells can be computationally ordered along ‘pseudotime’ and changes in the expression profiles of individual genes can be subsequently determined. However, while cell fate decisions are typically associated with profound changes in expression, many such changes are downstream of the initial cell fate decision. Instead, subtle changes in patterns of variation and coexpression of genes across developmental time, sometimes not associated with substantial changes in mean expression, have been argued to play a more critical role in symmetry breaking^{6,7}. Consistent with this, higher order interactions (i.e., looking beyond changes in mean expression) have proved highly informative for understanding genomics data, for example in supervised machine learning settings⁸ and for estimation of unknown spatial patterning⁹. Additionally, with recent developments in high-throughput and high-resolution spatially resolved gene expression mapping (e.g., Spatial Transcriptomics¹⁰; seqFISH¹¹; MERFISH¹²) it is now possible to explore the relationship between higher-order interactions and spatial location. For example, in the context of embryogenesis, do small numbers of spatially-localized cells display aberrantly higher variability in expression profiles prior to committing to a downstream fate?

From a computational perspective, methods for studying higher-order interactions are currently lacking. Although numerous methods have been developed for ordering cells along pseudotime, a computationally derived prediction of cell-type differentiation trajectories^{13–16}, methods for identifying individual genes that significantly change their expression levels across the pseudotemporal trajectory^{17–19} typically focus on changes in mean expression of single genes and do not characterize subtle changes in patterns of covariation between subsets of genes across this trajectory. In those cases where higher-order interactions have been studied, a typical analysis aiming to compare correlation patterns along pseudotime first defines strict nonoverlapping sets of cells before estimation of a covariance network, either through direct thresholding on the correlation matrix or using other methods^{20,21}. However, estimation of such networks is noisy²², and ignores potentially subtle but consistent changes across a continuum, as well as requiring an often ad hoc dichotomization or classification of cells into discrete groups. Even in the situation where the data arises from two distinct samples with unknown labels, we have previously shown using simulations that two-sample differential correlation methods are not particularly robust under model mis-specification²³. As we have also previously discussed²³, treating the

sample ranking as a covariate and testing for an interaction effect in a linear model is restricted to identifying linear and thus monotonic interactions, which may not be present, especially in highly dynamic or complex trajectories going through multiple changes in the differentiation process. In the context of spatially resolved gene expression data, fewer methods exist, with the focus being on testing the existence of pre-defined patterns²⁴ (e.g., a signaling gradient); however, these require *a priori* knowledge about the spatial structures of interest.

Results

Single cell Higher Order Testing (scHOT)

Here we introduce single cell Higher Order Testing (scHOT), a framework for examining changes in higher order structure, such as correlation among genes across differentiation pseudotime, among discrete groups, and across spatial landscapes. scHOT builds on our previous work, DCARS (Differential Correlation Across Ranked Samples), which used bulk RNA-sequencing data to test for changes in gene-gene correlation across ranked individual samples²³. Our approach requires one of the following types of cell-specific information (Figure 1): A) a ranking of cells, which will typically be across pseudotime, or B) spatial coordinates in either two or three dimensions. In the case where spatial coordinates are inferred²⁵, scHOT is also applicable using either the cell ranking along a gradient or in the inferred 2/3D space. Given this cell-specific information, as well as a scheme for determining local sample-specific weights, we calculate local higher order interaction vectors among single genes or pairs of genes, uncovering local changes in variability or covariation respectively (Figure 1). Sample-wise permutation testing is then used to assess statistical significance, while retaining the global variability or correlation structure of the original data. This framing of the significant genes and gene-pairs in terms of the set of local higher order interaction estimates allows patterns of changes across the trajectory, groups, or space to be characterized in terms of the higher order interaction, rather than simply by changes in the mean. Moreover, scHOT identifies groups of genes for which similar higher order patterns arise. For a more detailed discussion of how scHOT can be applied in practice, see the Supplementary Note.

scHOT identifies multiple higher order associations during liver development

We first analyzed four single-cell RNA-sequencing datasets designed to study the early development of the mouse liver^{26–29} (Methods, of which three contained hepatic cells). The integrated data encompassed 7 days of development (from embryonic (E) day 10.5 to E17.5), which covers the period where progenitor hepatoblasts transition towards more mature hepatocytes and cholangiocytes (Figure 2A). As expected, when using Monocle 2¹⁹ to order the cells in pseudotime, we observed a clear bifurcation where hepatoblasts differentiated into either cholangiocytes or hepatocytes (Figure 2B). This was supported by a higher proportion of differentiated cells in the later embryonic time points as well as cell-type specific expression of known marker genes^{28,30,31} (Figure 2C).

Building on this, we first examined higher order patterns as cells transitioned from naïve hepatoblasts towards the bifurcation point where they commit to one of the two downstream

lineages (Figure 2D). In total 68 genes showed a change in variability (false discovery rate (FDR) adjusted P-value < 0.1 , all shown in Extended Data Figure 1A and 1B and Supplementary Table 1A) along the trajectory in the absence of changes in gene expression. Of these, 58 (85%) displayed significantly increased variability along the branch (Figure 2D). These genes were enriched for involvement in processes associated with cell division, chromosomal organization and DNA replication (Figure 2E; Methods), consistent with the notion that increased plasticity can precede cell fate commitment^{6,7}.

We next focused on the full trajectory from naïve hepatoblast through to hepatocytes. Specifically, we investigated whether scHOT could identify changes in correlation, in the absence of differential expression, thus providing insight into the potentially complementary set of gene regulatory modules that are activated during the process of commitment from hepatoblast to the hepatocyte lineage. Correlation patterns identified as cells transition from naïve hepatoblast to cholangiocytes can be found in Extended Data Figure 2A and 2B.

When focusing on the hepatoblast to hepatocyte lineage, we identified numerous changes in correlation between pairs of genes that did not change their individual mean expression (Methods). An example of such a gene-pair is *Cdt1* and *Top2a* (Figure 3A, FDR adjusted P-value < 0.03), which are protein-protein interacting partners³² that have been implicated in regulation of the cell cycle in human and mouse stem cells³³. This pair of genes changes from being strongly negatively correlated in the progenitor population to displaying no correlation in the more differentiated hepatocytes. Interestingly, when considering each gene separately, neither *Cdt1* nor *Top2a* are significantly differentially expressed along the trajectory, or significantly differentially variable (FDR adjusted P-value = 0.70 and 0.12 respectively), indicating that the association between these two genes would not be identified without using scHOT. *Top2a* encodes a DNA topoisomerase, which controls and alters the topologic states of DNA during transcription³⁴, while *Cdt1* is a chromatin licensing and DNA replication factor that is required for DNA replication and mitosis³⁵. Our observation that these genes move from being negatively correlated to displaying no correlation suggests a trade-off between chromatin remodeling and transcription at the earlier stages of differentiation, potentially facilitating both proliferation and the global changes in gene regulatory architecture that arise when cells commit towards the hepatoblast lineage.

Across all 22,155 gene-pairs tested, 224 displayed different patterns of correlation (FDR adjusted P-value < 0.2), encompassing 136 unique genes (Supplementary Table 1B). Gene-pairs that were differentially correlated were not found to be associated with genes that were also differentially variable along the trajectory (Fisher's Exact Test P-value > 0.4) for either hepatocyte or cholangiocyte branch, suggesting an independent relationship between changes in correlation of gene-pairs and variability of the genes. The majority of local correlation patterns of these gene-pairs exhibited either a 'gain' or a 'loss' of correlation, across developmental time reflecting the prior understanding of a continuous differentiation towards the end fate (Figure 3B). Using the local correlation patterns as input, we performed hierarchical clustering to group these gene-pairs into 9 clusters. Functionally annotating the genes belonging to these clusters revealed that, in general, clusters associated with a loss of correlation (e.g., Cluster 8) contained genes linked with DNA replication and cell division.

By contrast, clusters that ‘gained’ correlation (e.g., Cluster 7) along the trajectory were associated with hepatocyte-linked functions such as lipoprotein particle remodeling, lipoprotein metabolism, as well as mitotic cell cycle and cell division (Figure 3C, all clusters shown in Extended Data Figure 2C). A small number of gene-pairs displayed more unexpected correlation patterns, with a transient peak of co-expression at an intermediate point along the trajectory (e.g., Clusters 5 and 9) near the bifurcation point, suggesting a transient role in cell fate commitment.

Finally, we explored whether the differential patterns of higher order interactions could reveal genes that were related specifically to differentiation into the hepatocyte or cholangiocyte lineages, or if they reflected a common pattern of exit out of the hepatoblast state into mature cells. To do this, we used network strength as a test statistic to characterize genes most related to a specific branch (hepatocyte or cholangiocyte), or common to both branches. Permuting gene labels over the topology of the gene network allowed assessment of statistical significance (Methods), revealing five genes that were shared among both branches, and ten and four genes significantly specific to the hepatocyte and cholangiocyte branches respectively (Figure 3D). In particular, the gene *Cdt1* appears associated to differentiation in general, i.e. from the hepatoblast state to either differentiated state, rather than any of the two terminal states. By contrast, *Apom* and *Apoa2*, encoding apolipoprotein, were more associated with hepatocyte function. More surprisingly, we identified the histone gene *H2afz* as more specific to the hepatocyte lineage, indicating a potential association with changes in global chromatin organization as cells commit towards a hepatocyte fate.

scHOT identifies local patterns of correlation in the mouse olfactory bulb

Finally, we considered whether scHOT could also be used to identify cryptic local correlation when gene expression information is available in a spatial context. Specifically, to date, most studies of spatially resolved gene expression data have focused on clustering cells into groups or testing known patterns of correlation – we reasoned that scHOT would provide an unbiased approach for identifying local patterns of correlation that might be missed by such approaches, which rely on changes in mean expression only. To this end we considered the mouse olfactory bulb (MOB), which displays a highly stereotypical structure, with clear patterns of concentric layers corresponding to granule, internal plexiform, mitral, external plexiform, glomerular and olfactory nerve layers moving from the inside out (Figure 4A), along with distinct patterns of gene expression along this space³⁶. Recently, spatial transcriptomics¹⁰ was used to measure gene expression levels in small spatially-distinct regions of the MOB, thereby facilitating the unbiased identification of patterns of gene expression in space.

Using scHOT, we identified a set of 167 gene-pairs as significantly differentially correlated across space (FDR adjusted P-value < 0.2), with 42 non-differentially expressed highly variable genes appearing at least once among this set (Methods). Interestingly, we found that numerous pairs of genes displayed diffuse patterns of local correlation that were not apparent when visually comparing their individual expression profiles. For example, *Arrb1* (beta-arrestin-1) is widely expressed in the brain³⁷, consistent with its diffuse expression across the MOB in the spatial transcriptomic data. Similarly, *mTOR* (mammalian target of

rapamycin), *Uchl1* (Ubiquitin carboxyl-terminal esterase L1), and *Dnm3* (Dynamin-3) are all broadly expressed^{38–40}. Nevertheless, we identified all three genes as positively spatially correlated with *Arrb1* (Figure 4B and Extended Data Figure 3A). Consistent with this correlation, *Arrb1* can regulate *mTor* activation⁴¹ and interact with the ubiquitination pathway to down regulate receptor signaling⁴². Further, both *Arrb1* and *Dnm3* function in endocytosis⁴³.

To explore more general patterns, we clustered all significant gene-pairs using their local spatial correlation patterns into 8 distinct groups (Figure 4C). Despite the relatively low-resolution of the data (spatial transcriptomic data is limited to a resolution of ~10 cells (approximately 100 μm)¹⁰), a variety of local correlation patterns were observed, often associated with distinct biologically meaningful regions of the bulb. Giving confidence in our analysis approach, clustering the cells based on their local correlation pattern largely recapitulated the symmetry of the MOB, with multiple cell groups corresponding to symmetric sets of cells, e.g. Cell groups 1, 4, and to some extent Cell groups 5 and 6 (Figure 4C). Additionally, Cluster 1 contained genes that were positively correlated within the Olfactory Nerve Layer and Clusters 4 and 5 were associated with the Mitral and External plexiform layers (Figure 4D). Functional annotation of the genes belonging to these clusters revealed associations with distinct neuronal terms including signaling events such as endocytosis (all clusters shown in Extended Data Figure 3B). Interestingly, “myelin sheath” was highly ranked in multiple clusters (Clusters 2, 4-7). In these clusters, the strongest patterns of correlated spatial expression occurred within more internal layers of the bulb, overlapping with the mitral and granule layers. This is consistent with the myelination of the lateral olfactory tract as it exits the bulb⁴⁴. Clusters 1, 3, and to some extent 8 and 9, in contrast, possess spatial correlation patterns that encompass more external layers such as the olfactory nerve layer, and genes within these clusters are not highly associated with the term ‘myelin sheath.’ This is consistent with the fact that olfactory sensory neurons entering the bulb in the more external olfactory nerve layer are not myelinated. In sum, we have shown here that exploiting higher-order structure can reveal unexpected and spatially-coherent regions of structured heterogeneity that persist in the absence of mean expression changes, and that approaches that focus only on the latter will fail to fully exploit the wealth of information contained within such data.

Discussion

In this paper we have demonstrated the utility of higher order testing for single cell data. We examined scHOT in the context of two biological systems with distinct data characteristics – liver development and the mouse olfactory bulb. scHOT is robust due to the choice of underlying higher order metrics such as rank-based Spearman correlation; powerful as it uses a permutation framework retaining the global variability and covariance structure for inference; and extremely flexible as it can be tuned by 1) Varying the local weighting scheme in terms of shape (triangular, block, any other user defined weight) and span, 2) Choice of underlying higher order effect function (weighted Pearson correlation, weighted Spearman correlation, weighted zero-inflated Spearman or Kendall correlation, or any choice of higher order estimate when using the block weighting scheme), and 3) Choice of summarization estimate for the local higher order vector (by default the standard deviation).

In general, this contrasts with other methods that estimate changes in expression across either a pseudotime trajectory or across space, which require a set of candidate hypotheses to test explicitly. In the spatial context, scHOT differs substantially to other methods such as SpatialDE²⁴, in that we can test either a single gene (identifying spatially variable) or two genes (identify spatially differentially correlated), and no prior suite of potential spatial structures need be provided to identify genes that are of interest.

From a biological perspective, the concept of characterizing coordinated changes over time could enable better characterization of the molecular processes underpinning cell fate choice. In particular, it will help us to better understand whether increased plasticity, as manifested by increased cell-to-cell variability, is a general feature that precedes all instances of cell fate commitment, or whether this is restricted to a specific set of early biological differentiation events in systems such as embryonic hepatoblast differentiation as found here, and other systems^{7,45–50}. Such heterogeneity could also be a driver of differential cell fate or cell function in a spatial context: specific patterns of local correlation could indicate that a specific region of a tissue or organ is primed towards a specific fate. Intriguingly, our reanalysis of data from the mouse olfactory bulb identified patterns that were observed in the absence of changes in mean expression but associated with known spatial structure of the bulb.

In summary, scHOT is a method for inference of changes in higher order interactions, not just changes associated with the mean, and as such offers a new lens to interrogate single cell data and describe patterns of variation and covariation, offering additional and complementary insight to that obtained by examining changes in mean expression. It is enabled by a statistical framework that captures nonlinear changes in variability and correlation structure by using sample ranking approaches to avoid having to discretize responses and risk obscuring biologically meaningful results. This is especially important for continuous single cell trajectories and for studying spatial structure within ostensibly homogeneous cell types. By facilitating such analysis, scHOT will enable investigations into how highly localized patterns of variation and co-variation influence cell fate and cell function.

Online Methods

Datasets

The following datasets were used to examine scHOT, and demonstrate its utility in extracting insights from diverse sources of single cell and/or spatially resolved data.

Developmental Liver Data—The ‘Developmental Liver Data’ is a full-transcript scRNA-Seq dataset generated using plate-based protocols from four distinct sources^{26–29}, across multiple mouse embryonic time points from Embryonic Day (E)10.5 to E17.5. The data were originally processed as size-factor standardized logCPM values per dataset, and integrated using scMerge⁵¹, taking advantage of genes that are found to be stably expressed in single cells⁵². These data comprise several cell types including hepatic cells such as hepatoblasts, cholangiocytes, and hepatocytes, among other cell types such as immune cells (Figure 2A). Monocle 2 was used to infer a differentiation trajectory exclusively for the

hepatic cells. We applied scHOT to these data, considering the following testing scenarios: changes in variability across the first branch from hepatoblasts to the cell fate decision point; and changes in correlation between pairs of highly variable genes along the entire branch from hepatoblasts to hepatocytes and the full branch from hepatoblasts to cholangiocytes.

To select genes for downstream analysis, we considered genes that were highly variable (HVGs)⁵³ (Extended Data Figure 4A) but that had consistent mean expression along the trajectory. To do this, we performed liberal differential gene expression testing along pseudotime by fitting, for each gene, two linear models (slope and intercept, and polynomial of degree two) and identifying a gene as differentially expressed if it was significant (F-test; unadjusted P-value < 0.05) in at least one of the tests when compared to an intercept only model. This differential expression testing was performed for the hepatoblast to hepatocyte trajectory and for the hepatoblast to cholangiocyte trajectory. The resulting sets of genes (i.e., highly variable and non-differentially expressed) were combined to form all pairwise combinations as the scaffold for higher order gene-pair testing.

Spatial data from the mouse olfactory bulb (MOB)—The ‘Mouse Olfactory Bulb’ data is a Spatial Transcriptomics dataset, where an array spotted with probes that have barcodes corresponding to defined locations was used to measure spatially-resolved gene expression levels¹⁰. We consider data where this technology was used to measure expression levels across a section of the mouse olfactory bulb (MOB), where each spatially resolved region contains a measure of gene expression averages across approximately 10 cells¹⁰. This cross section of the MOB comprises concentric layers visible with H&E staining (Figure 4A), associated with the granule, internal plexiform, mitral, external plexiform, glomerular and olfactory nerve layers moving from the inside out. The resulting expression data is derived from high throughput RNA sequencing using barcodes corresponding to the spatial locations, as well as unique molecular identifier (UMI) barcodes. We preprocessed these data to obtain size-factor standardized log-transformed counts for each gene and spatial sample. We identified genes as spatially differentially expressed by performing scHOT using a first-order metric of local weighted mean expression (2,542 differentially expressed genes; unadjusted P-value < 0.05). After identifying the intersection between genes that were highly variable⁵³ (Extended Data Figure 4B) but not differentially expressed we used scHOT to test all pairwise combinations for this set of genes.

Choice of local weighting scheme

For the trajectory-based analyses we selected a triangular weight matrix with a span of 0.25. For the spatial-based analysis we selected a two-dimensional triangular weight matrix (i.e. a cone) also with a Euclidean distance span of 0.05 (here corresponding roughly to 9 surrounding sampled points).

scHOT test statistic and inference

For single gene testing we use a local weighted variance estimate

$$S_j^2 = \sum w_{ij}^* \sum w_{ij}^* x_i^2 - (\sum w_{ij}^* x_i)^2,$$

$$w_{ij}^* = \frac{w_{ij}}{\sum_i w_{ij}},$$

where w_{ij} is the cell-specific weight for cell i and position j , and x_i is the gene expression measure for gene x and cell i , and all summation is performed over index i . For testing pairs of genes we use a weighted Spearman correlation

$$c_j = \frac{\sum w_{ij} \sum w_{ij} r_{xij} r_{yij} - \sum w_{ij} r_{xij} \sum w_{ij} r_{yij}}{\sqrt{(\sum w_{ij} \sum w_{ij} r_{xij}^2 - (\sum w_{ij} r_{xij})^2)(\sum w_{ij} \sum w_{ij} r_{yij}^2 - (\sum w_{ij} r_{yij})^2)}}$$

where

$$r_{xij} = \text{Rank}(x_i \times I\{w_{ij} > 0\}),$$

and

$$r_{yij} = \text{Rank}(y_i \times I\{w_{ij} > 0\}),$$

where additionally y_i is the gene expression measure for gene y and cell i , and all summation is performed over index i .

The scHOT test statistic is a measure of variability associated with this vector of local variances or correlations. To compute this, we first calculate the sample standard deviation to estimate the variability associated with the set of local variance estimates $\{s_j^2, j = 1, 2, 3, \dots\}$ or local correlation estimates $\{c_j, j = 1, 2, 3, \dots\}$.

Statistical testing is then performed by randomly permuting cell labels, while keeping the overall gene expression structure constant. Thus, within each permutation round, the global correlation or global variance remains the same, while the vectors of local variability or local correlation vary. In all cases, we used sample or cell permutation and defined significance by controlling for a 0.2 Benjamini-Hochberg⁵⁴ FDR in all differential correlation tests, and at 0.1 FDR for variability tests. For correlation-based tests, we used the fact that the null distribution is based only on the two matched gene expression vectors to interpolate null distributions given the global correlation value and the number of samples (see Computational efficiency section) in order to speed up computation. In the case of discrete groupings, we use a normal approximation to the null distribution to estimate high resolution P-values.

Downstream analysis

After identifying significant gene-pairs, we took their local correlation profiles and hierarchically clustered them to identify patterns of differential correlation across either

pseudotime or space, using maximum distance and complete linkage hierarchical clustering. Maximum distance was used in this case since we wish to capture similarity of profile shape as well as absolute distance. For the Liver Developmental Dataset correlation analysis, we smoothed local correlation vectors before hierarchical clustering using loess. For both datasets we extracted discrete clusters from the hierarchical clustering using the R function `cutree` with number of clusters estimated using `dynamicTreeCut`⁵⁵. To functionally annotate these clusters, we performed gene set enrichment analysis using mouse Gene Ontology terms with between 10 and 500 genes appearing in each dataset, and at least one gene appearing from the testing scaffold^{56,57}, using Fisher's Exact Test to test for overrepresentation of genes, using all scHOT tested genes as the gene universe. An FDR adjusted P-value < 0.05 was considered to be statistically significant.

Comparing between trajectory branches

We implemented a statistical test for comparing the change in correlation between the two branches, by examining the normalized network strength across the tested networks per branch. We defined network strength for a given gene (node) as the sum of edge weights for significant gene-pairs associated with the gene, divided by the total edge weights across the entire network. The edge weight we selected was the $-\log(\text{FDR adjusted P-value})$ for each gene-pair. For each gene, we calculated the network strength of all genes per branch. We then compared these network strength values between branches using an MA-plot, i.e. comparing the sum of network strengths with the difference of network strengths. To assess significance associated with a single gene – i.e. a gene that tends to have a higher network strength than expected by chance, we repeatedly permuted the gene-pair edge weights across the network and calculated the permuted MA-plot. Individual genes were identified with a significantly nonzero network strength using the Euclidean distance from the origin as the test statistic. To identify genes with a branch-specific network strength, we considered the ratio of significance towards each branch as the test statistic.

Computational efficiency

We previously observed a relationship between the total number of samples and the null distribution of the DCARS test statistics²³. Here we uncovered further association between the null distribution of the scHOT test statistics and the global correlation across all samples. This represents an opportunity to significantly decrease computational time as one can 'borrow' permutations from similarly distributed genes and gene-pairs to estimate the P-value. Our approach is to first calculate global correlations for all gene-pairs to be tested, and then take a uniform sample among the gene-pairs according to the global correlations. For this subset of gene-pairs we permute sample labels and calculate scHOT test statistics. Then for any given gene-pair of interest, we extract the desired number of permutations from this set of permuted scHOT test statistics, according to how similar their global correlation is. These are shown in Extended Data Figures 4C and 4D for the liver and MOB data respectively. To examine the accuracy of estimated P-values, we performed the full 10,000 permutation testing for all gene-pairs for the liver hepatocyte branch, artificially setting the zero P-values to 1/10,001, and observed a very high concordance between calculated and estimated P-values, with a Spearman correlation of 0.995 (Extended Data Figure 4E). We

observed a slight uptick in the loess local smoothed curve fit, reflecting the fact that we were able to use more than 10,000 permutations for estimation due to the borrowing procedure.

schHOT extensions and considerations

schHOT is a flexible framework within which multiple aspects can be modified to facilitate bespoke analysis. For example, higher order patterns can be studied along trajectories, across space, or among discrete groups (Extended Data Figure 5A). Moreover, distinct sets of genes or gene-pairs can be interrogated depending on the biological question of interest (Extended Data Figure 5B). Of particular interest, the local weighting scheme and concordance function can also be adapted, depending upon the biological context (some examples are given in Extended Data Figure 5C). For example, if one were interested in identifying changes in higher order interactions along a circular trajectory, e.g. the cell cycle, one could define a local weighting scheme that was also circular – by ensuring that the two ends match given any starting point. Another example is for discrete groups that are either completely distinct, or ordered in some way – e.g. over discrete time points along a time course experiment, one may define the weight matrix to incorporate the discrete grouping, while also accounting for the flanking groups. More generally, one may wish to place a higher local weight over a particular local region and a smaller weight over a different region. We note, however, that these changes to the weighting scheme may affect the generalizability of the null distributions across multiple genes or gene-pairs, so the user should take care in ensuring that the null distributions appear similar when employing computational speed-up steps.

Any concordance metric can be ‘plugged in’ if using a binary weighting scheme, representing a ‘block’ type of weight matrix. That is, one is able to use the fast implementation of distance correlation between two distributions⁵⁸, mutual information, partial correlation, or any concordance metric suited especially to other data types such as ordinal or binarized single cell data⁵⁹, without needing to explicitly derive weighted formulations of these concordance metrics. Additionally, any concordance metric that doesn’t necessarily have a ‘weighted’ formulation and/or implementation can be utilized using the block weighting scheme. This makes schHOT versatile, by enabling user-defined metrics. For summarizing the vector of local higher order statistics, users may wish to substitute the sample standard deviation with any other choice of variability or change estimate – e.g. if the goal is to examine how monotonic a change in higher order interaction is, a measure of monotonicity such as mutual information or Spearman correlation with the weighting scheme index could be used.

Single cell RNA-Seq data can have vastly different statistical structure and properties depending on the technology used, for example using plate-based or droplet-based technologies. Here, a particular issue to consider is the sparsity of gene expression measurements, since this can impact on reliable estimation of higher order statistics like variability or correlation. In general, robust statistical measures, such as Spearman correlation or median absolute deviation (MAD) are worth considering in the presence of outliers (or inliers). Other bespoke metrics such as the zero-inflated Kendall’s tau⁶⁰ may be worthwhile, but many observations would be needed to accurately estimate these higher

order statistics. Span may be chosen with a larger or smaller value depending on the sparsity of the data, with span increasing as sparsity (and often sample size in terms of number of cells) increases.

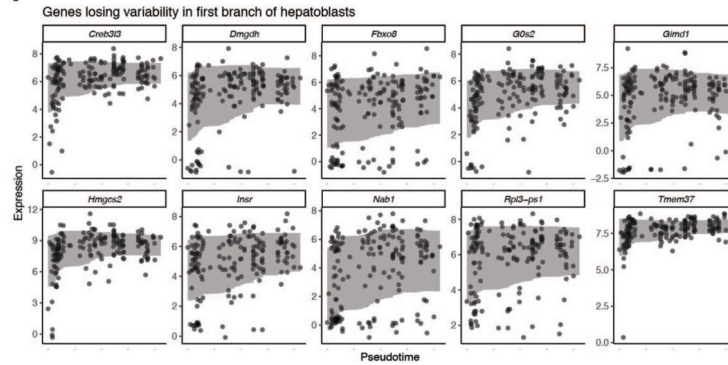
To understand how the choice of span influenced the results, we performed differential correlation testing using a range of spans for both the liver hepatocyte trajectory and for the spatial MOB data. We found that scHOT results were robust over small variations in span choice, with high concordance in P-values for both datasets (Extended Data Figure 6A-B and 5C-D). For the trajectory data, we noticed there was a slight difference in the gene-pairs identified as significant for small span (0.35 and below) and large span (0.40 and above). Closer inspection revealed that gene-pairs that were identified only with a small span tended to correspond to high degree of change in local correlation pattern, whereas those identified only with a large span tended to correspond to more subtle changes in local correlation patterns (Extended Data Figure 6E). This agrees with our intuition that a small span has enough resolution to capture local changes in higher order structure, while larger spans, which use more data points and thus yield better estimates of higher order statistics, may be more powerful in detecting subtle changes in higher order structure. Consequently, if changes in local higher order structure are likely to be monotonic, it is better to select a larger span value, but if capturing multiple changes in higher order structure along a long trajectory with multiple transient states is of interest, we suggest using a smaller span. Interestingly, in the spatial setting, we did not observe any increase in significant gene-pairs with increase in span. This may highlight the difference in data structure between cells ordered along a trajectory and cells in their spatial context, where a span corresponding to the immediate surroundings (here the span of 0.05 corresponds to around 13 neighboring points) is more informative. However, in a spatial context where there is likely to be a long-range gradient of higher order structure, using a larger span may detect more genes and sets of genes of interest. For typical use we suggest selecting a span of 0.25, if there are 400 or more cells; with smaller span values useful with higher number of cells; else use a proportion such that at least 30 cells are sampled at any given stage. In the spatial context we set a default of 0.05, and we suggest, for equally spaced transcriptomic readout, to select a span so that approximately 10-15 neighboring points are included.

To understand the effect of the choice of underlying metric for higher order testing, we examined how the sets of significant gene-pairs differed when alternative metrics were used. Specifically, we compared weighted Pearson correlation, weighted Spearman correlation, block-weighted Pearson correlation, block-weighted Spearman correlation, Brownian distance correlation (BDC), and maximal information criterion (MIC) as underlying metrics to use for scHOT. Note that since there is no obvious way to derive a weighted version of MIC and BDC we implemented them in a block-weighted manner, and thus also performed Pearson and Spearman in a block-weighted manner to link any differences to the metric itself and not the weighting scheme. We found that overall there was good concordance of results for all metrics with the exception of the MIC (Extended Data Figure 6F), which appeared to lack statistical power (Extended Data Figure 6G). More generally, we identified the highest number of significant gene-pairs using the Spearman correlation metric (Extended Data Figure 6G), with minimal difference between the triangular or block weighting schemes.

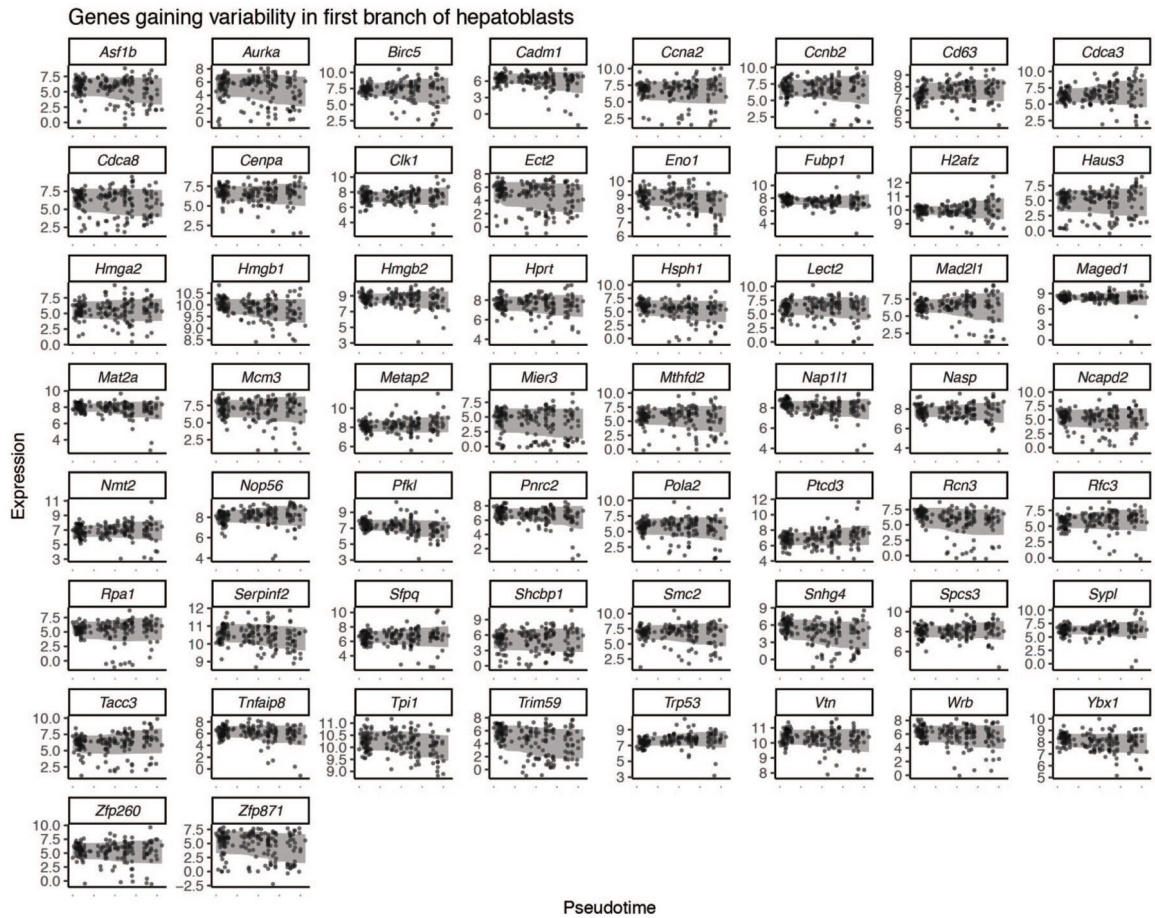
To explore the level of robustness and stability of scHOT testing under random subsampling, we took 90%, 80%, 70%, 60% and 50% random subsets - repeated 10 times - of the cells of the hepatoblast to hepatocyte branch and performed differential correlation testing with scHOT. We found that the P-values are highly correlated with the full dataset result, especially at the 90% threshold, with only a relatively small number of genes being excluded (Extended Data Figure 7A-B). We also note that as we decrease the sample size substantially, the P-value distribution trends towards more conservative than the full dataset, which is to be expected (Extended Data Figure 7C). We note that while performing these analyses we ensured the span was consistent between the full data and the subsampled data, by setting the span value as the product of the subset fraction and the default span choice of 0.25.

Extended Data

A

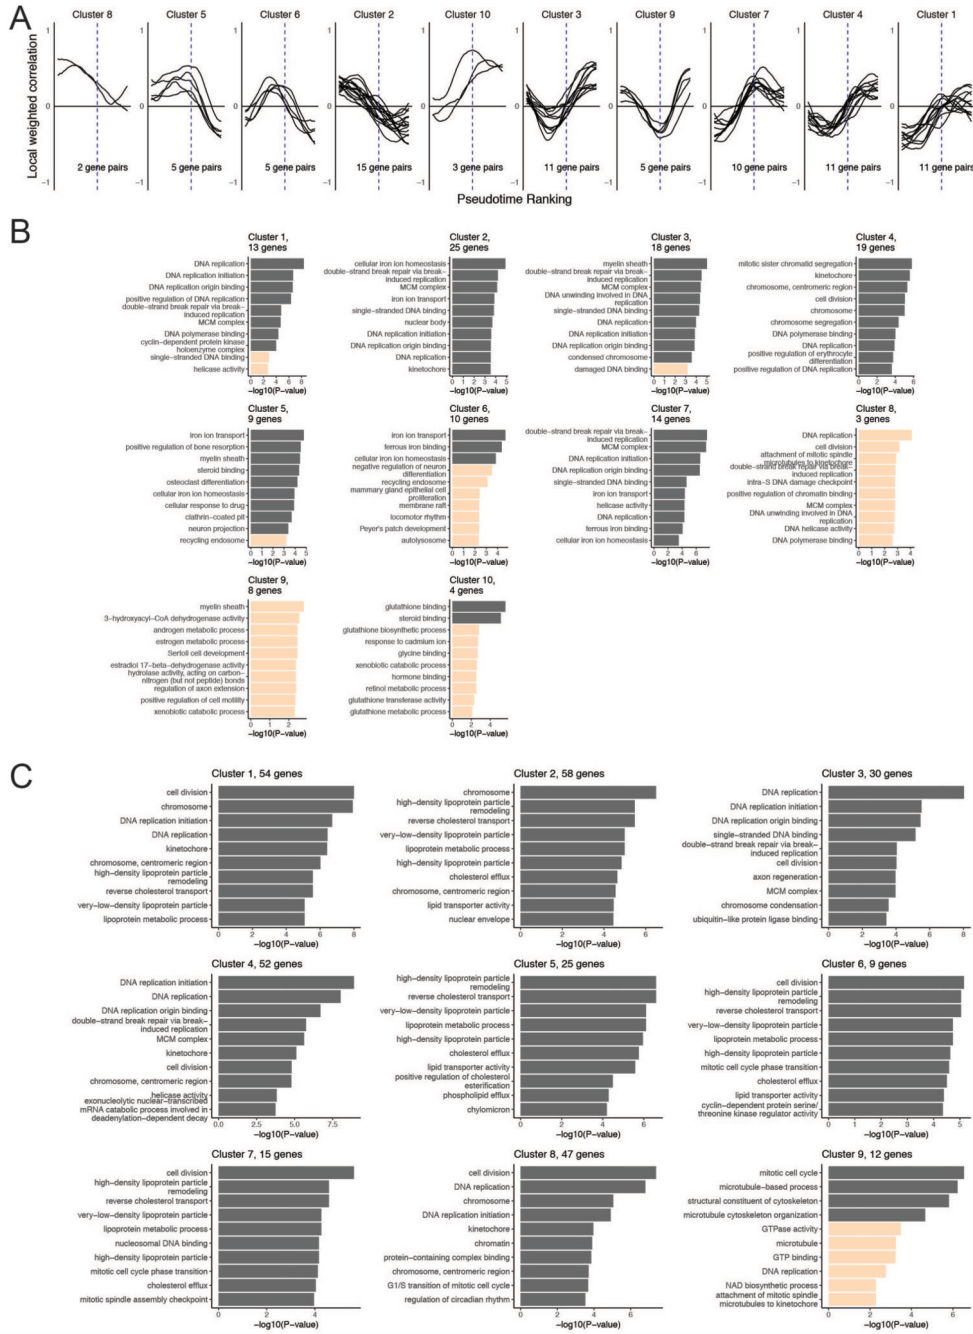


B



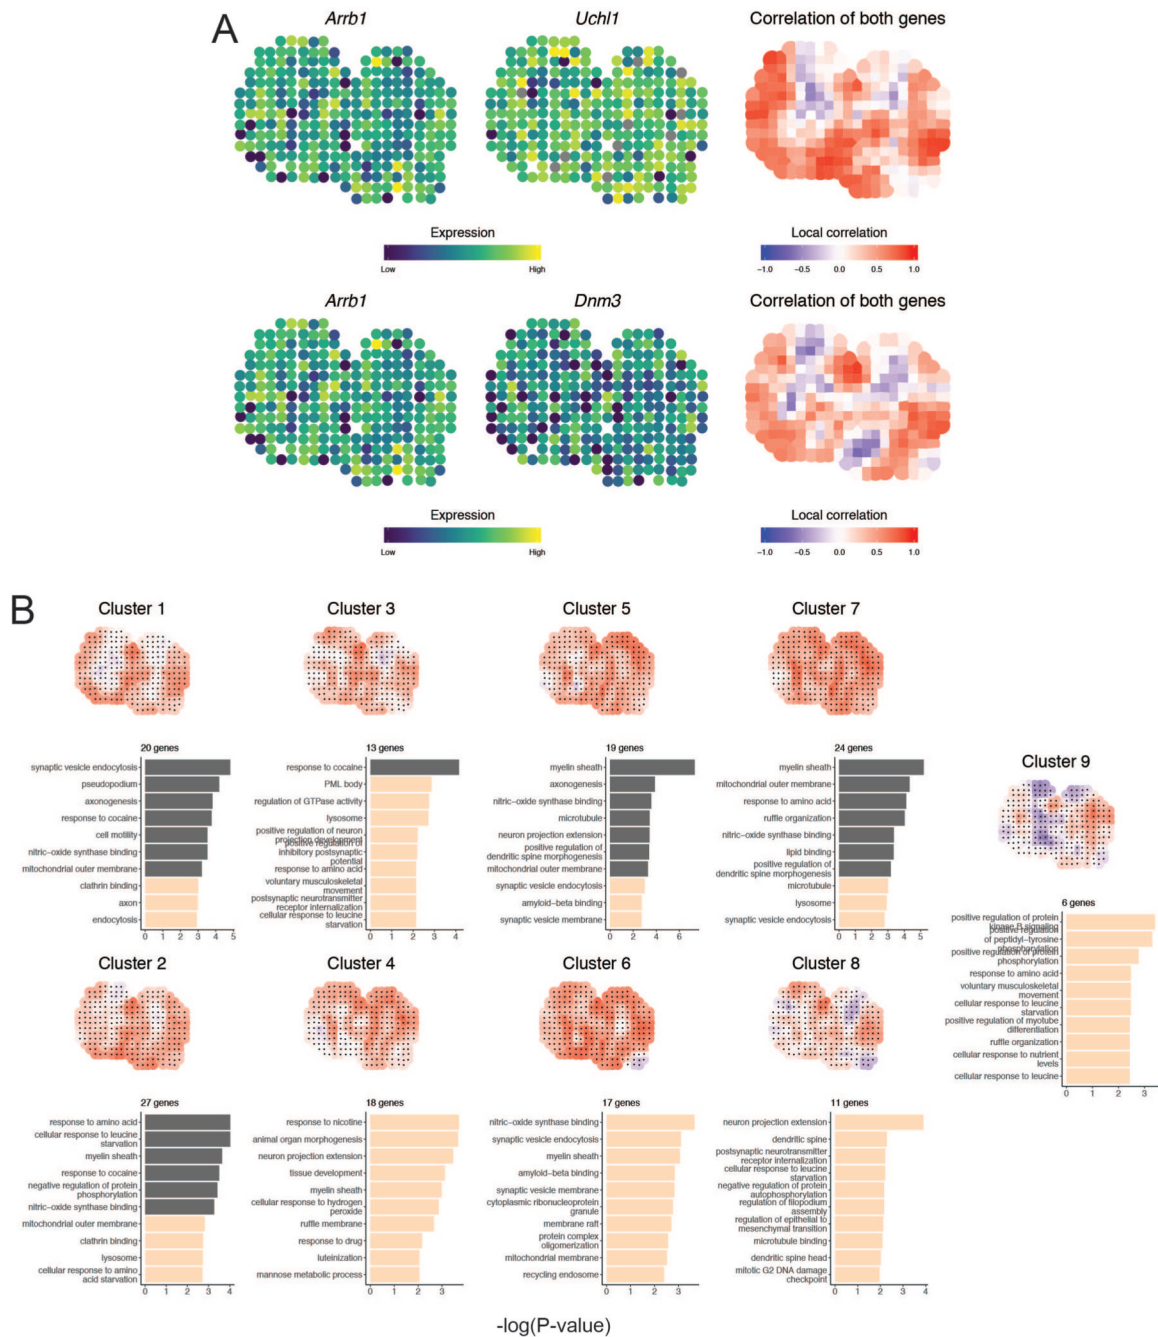
Extended Data Fig. 1. Variability testing along hepatoblast branch using scHOT.

A. Scatter and ribbon plot of all significant genes showing loss of variability along hepatoblast branch (n = 408), ribbon width corresponding to adding and subtracting the weighted standard deviation from the weighted mean. **B.** Scatter and ribbon plot of all significant genes showing gain of variability along hepatoblast branch, with ribbon width as in Panel A.



Extended Data Fig. 2. Correlation testing of cholangiocyte branch using scHOT.

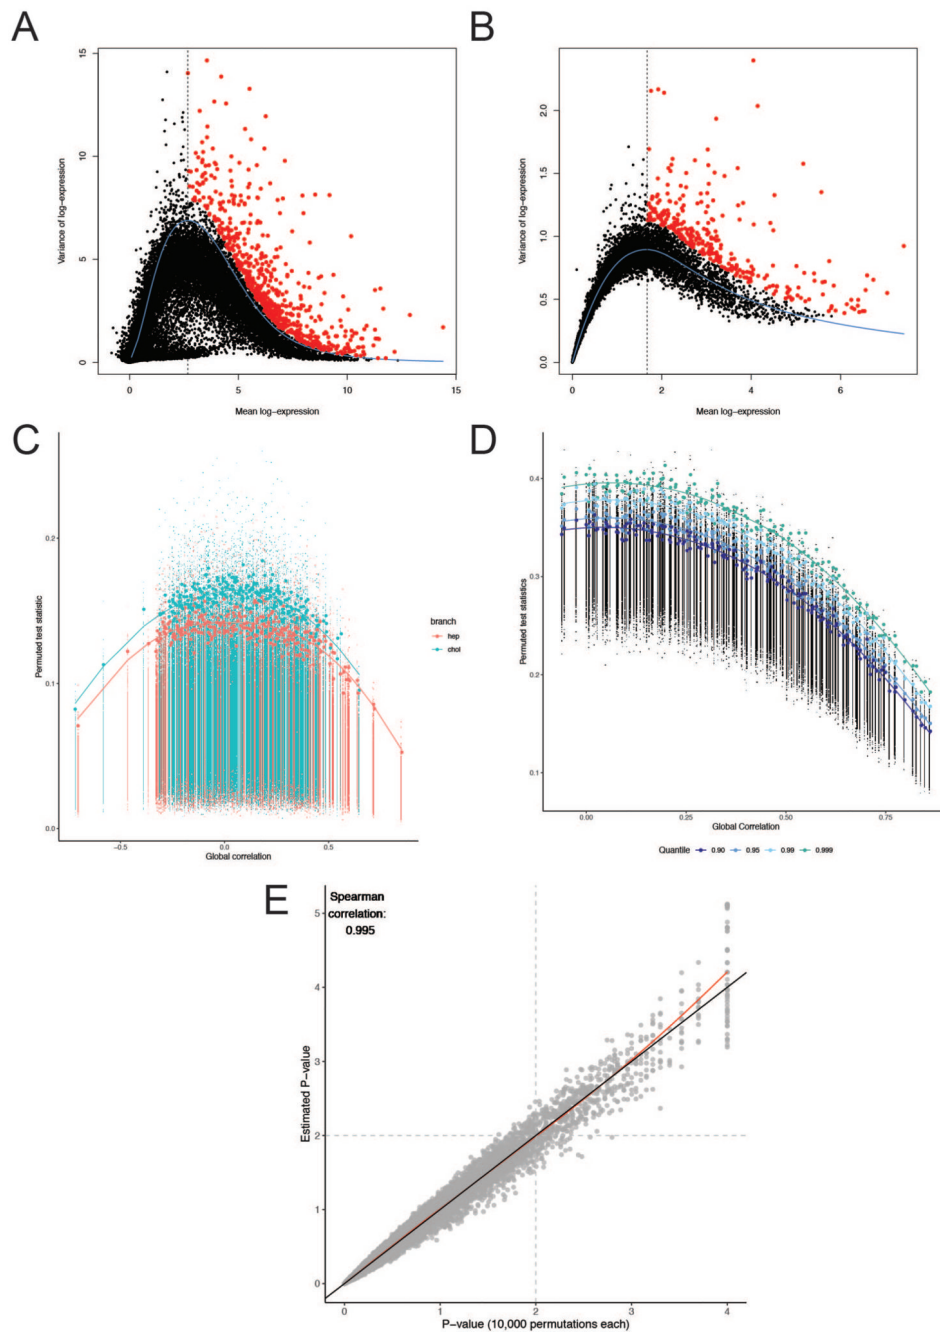
A. Line plots of clustered significant scHOT gene-pairs with FDR adjusted P-value < 0.2 for full cholangiocyte branch (n = 308). Vertical dashed line indicates trajectory branchpoint. **B.** Gene ontology functional enrichment using one-sided Fisher’s Exact Test barplots for all scHOT cholangiocyte clusters, grey bar color corresponds to FDR adjusted P-value < 0.05. Gene sample sizes indicated in each plot title. **C.** Gene ontology functional enrichment using one-sided Fisher’s Exact Test barplots for all scHOT hepatocyte clusters, with color as in panel B. Gene sample sizes indicated in each plot title.



Extended Data Fig. 3. Correlation testing of spatial data using scHOT.

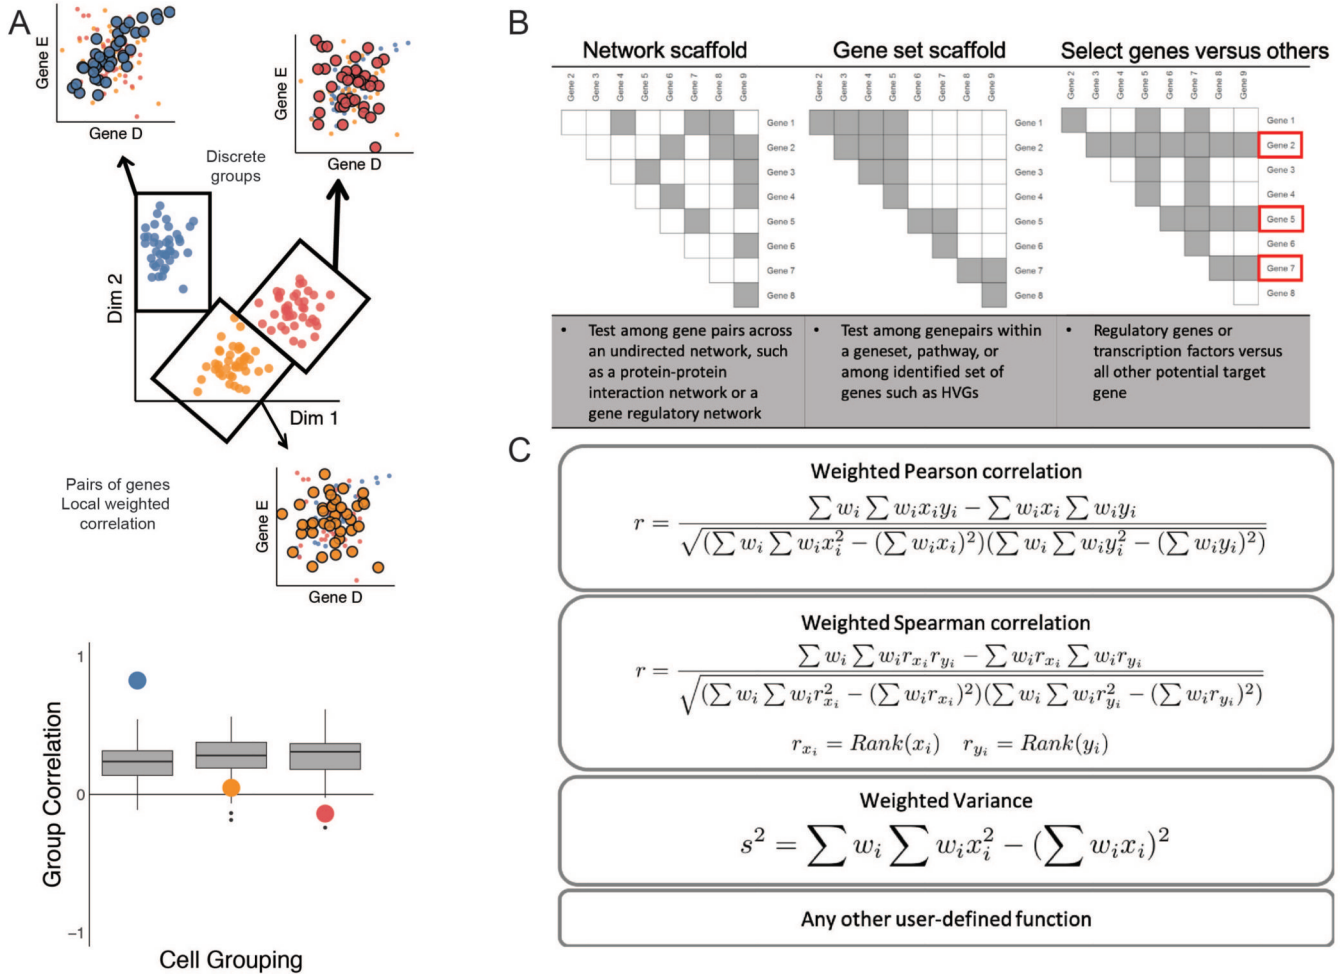
A. Spatial expression plots ($n = 262$ spatially resolved positions) of two gene-pairs *Arrb1* and *Uchl1* as well as *Arrb1* and *Dnm3* which are not significantly differentially expressed across space using scHOT, but are significantly differentially correlated across space using scHOT with FDR adjusted P-value < 0.2 . The third plot shows the local spatial correlation estimated for these two genes, recapitulating the layered pattern of the olfactory bulb. **B.** Spatial maps of mean local correlation and Gene Ontology functional enrichment using one-

sided Fisher's Exact Test barplots for all MOB scHOT clusters, grey bar color corresponds to FDR adjusted P-value < 0.05. Gene sample sizes indicated in each plot title.



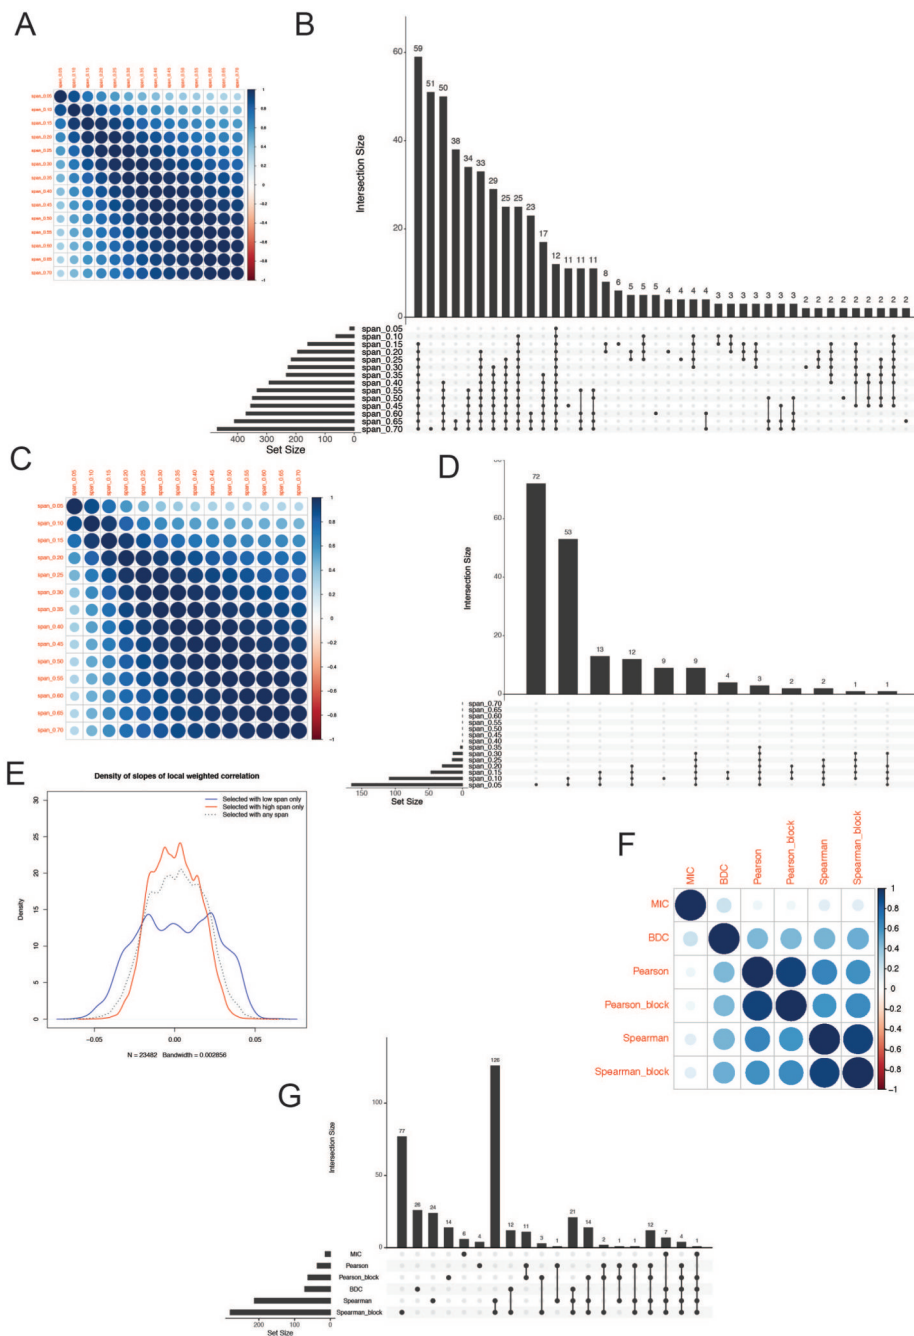
Extended Data Fig. 4. Highly variable gene selection and permutation testing for scHOT. **A.** HVG selection for Developmental Liver Data. **B.** HVG selection for Spatial Transcriptomics analysis. **C.** Global correlation and null scHOT correlation test statistics for sampled gene pairs in both hepatocyte and cholangiocyte branches ($n = 268$ and $n = 265$ respectively). **D.** Global correlation and null scHOT correlation test statistics for sampled

gene pairs (n = 172) for spatial MOB data. **E.** Scatterplot of $-\log_{10}(\text{P-values})$ for differential correlation testing (n = 22,155 gene pairs) of the liver hepatocyte branch, calculated using 10,000 permutations for each gene-pair (x-axis), and estimated using borrowed permutations over a subset of gene-pairs (y-axis). Black solid line corresponds to $y = x$, grey dashed lines correspond to unadjusted P-values of 0.05, and the solid red curve corresponds to the fitted loess curve.



Extended Data Fig. 5. Further scHOT method description.

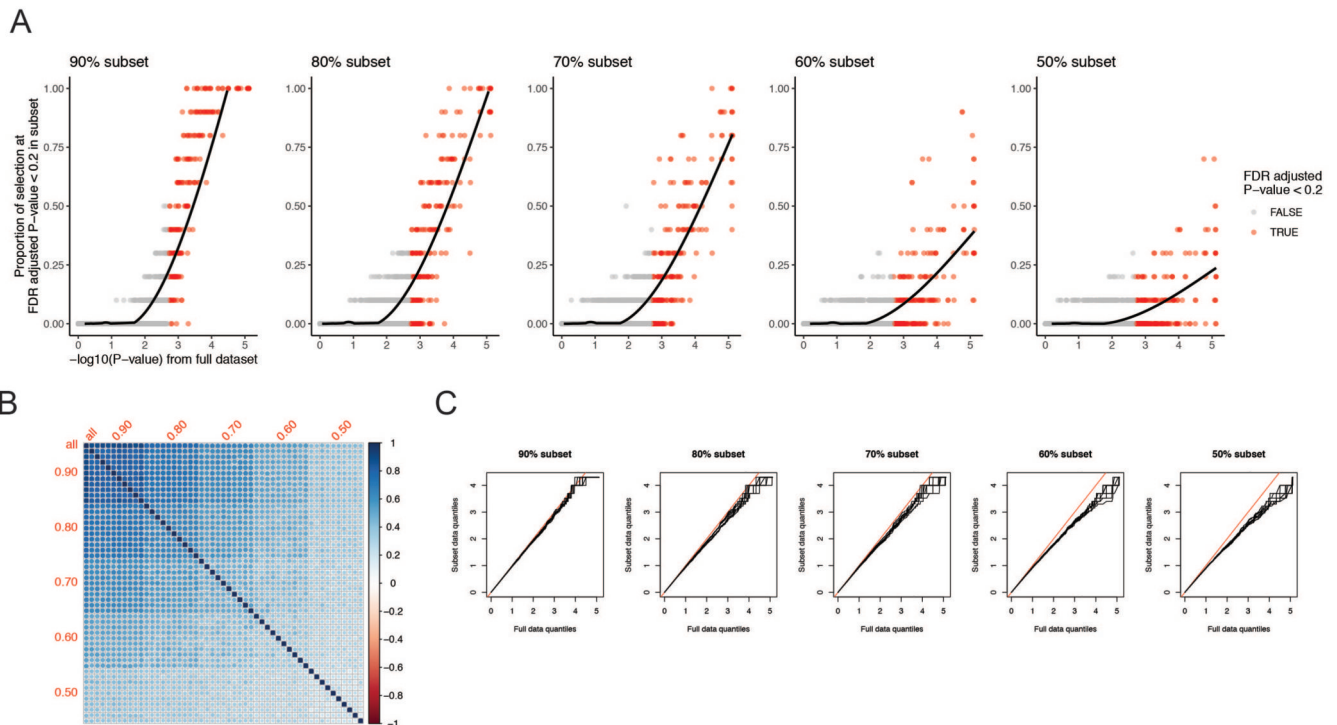
A. Illustrative example showing testing for correlation differences in three distinct groups. A set of local higher order statistics are calculated, and significance is compared by repeatedly permuting samples (grey boxplots). Illustrative example shows the set of local estimates of higher order statistics are combined using the sample standard deviation to assess how variable they are between groups. **B.** Possible schemes for the testing scaffold using gene networks, including: i) a gene-gene network; ii) a gene set scaffold where all pairwise combinations within a gene set are included; and iii) selected genes of interest versus all others. **C.** Examples of weighted higher order functions including weighted Pearson correlation, weighted Spearman correlation, weighted variance. Note that any user defined function can be used.



Extended Data Fig. 6. sCHOT stability under parameter choices.

A. Spearman correlation map of $-\log_{10}(\text{P-values})$ of sCHOT differential correlation testing ($n = 22,155$ gene pairs) in hepatocyte branch with different choices of triangular span, from 0.05 to 0.70 in steps of 0.05. **B.** UpSet plot of 509 significant gene-pairs (FDR adjusted P-value < 0.2) from each sCHOT testing scheme as in panel A. **C.** Spearman correlation map of $-\log_{10}(\text{P-values})$ of sCHOT differential correlation testing ($n = 903$ gene pairs) in MOB with different choice of spatial span. **D.** UpSet plot of 181 significant gene-pairs from each testing scheme as in panel C. **E.** Density plot of slopes of local correlation patterns of gene-

pairs selected with low span (0.35 and below, colored blue) and high span (0.35 and above, colored red), dotted line shows the density over all slopes. **F.** Spearman correlation map of $-\log_{10}(\text{P-values})$ of scHOT differential correlation testing ($n = 22,155$ gene pairs) in hepatocyte branch with different choices of higher order statistic, MIC – maximal information criterion, BDC – Brownian distance correlation, Pearson_block and Spearman_block refer to Pearson and Spearman correlation respectively, applied in a block-weighted context. **G.** UpSet plot of 367 significant gene-pairs from each testing scheme as in panel F.



Extended Data Fig. 7. Robustness of scHOT results.

A. Scatterplots of $-\log_{10}(\text{P-values})$ for hepatoblast to hepatocyte correlation scHOT testing ($n = 22,155$ gene pairs) against the inclusion frequency of gene-pairs with FDR adjusted P-values < 0.2 for repeated subsampling without replacement of 90%, 80%, 70%, 60%, and 50% of the cells from the trajectory. Red points correspond to those selected as FDR adjusted P-value < 0.2 criteria for the full dataset. **B.** Spearman correlation map of $-\log_{10}(\text{P-values})$ of each subsampling strategy. **C.** Quantile-quantile line plots of $-\log_{10}(\text{P-values})$ for the full data (x-axis) and $-\log_{10}(\text{P-values})$ for each subsampling scenario (y-axis), split by subsampling percentage. Red lines correspond to $y = x$.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

The authors thank all their colleagues, particularly at Cancer Research UK Cambridge Institute and The University of Sydney School of Mathematics and Statistics for their support and intellectual engagement. In particular, the authors gratefully acknowledge Pengyi Yang and Michael Morgan for their helpful discussion.

Funding

The following sources of funding are gratefully acknowledged. Royal Society Newton International Fellowship (NIFR1\181950) and funding from the Judith and Coffey Life Lab at the Charles Perkins Centre to S.G.. Australia NHMRC Career Developmental Fellowship (APP1111338) to J.Y.H.Y.. Research Training Program Tuition Fee Offset and Stipend Scholarship and Chen Family Research Scholarship to Y.L.. NIH grant (R21DC015107) to D.M.L.. SJTU-USYD Translate Medicine Fund-Systems Biomedicine (AF6260003) to J.Y.H.Y., X.S., and Z.G.H.. Core funding from EMBL and Cancer Research UK (award no. 17197) to J.C.M..

The funding source had no role in the study design; in the collection, analysis, and interpretation of data, in the writing of the manuscript, and in the decision to submit the manuscript for publication.

Data Availability

All data analysis was performed on publicly available data. The liver developmental dataset and description is available from https://sydneybiox.github.io/scMerge/articles/Mouse_Liver_Data/Mouse_Liver_Data.html and the specific R data file downloaded from www.maths.usyd.edu.au/u/yingxinl/wwwnb/scMergeData/liver_scMerge.rds

The mouse olfactory bulb data was downloaded from the Spatial Research website <https://www.spatialresearch.org/resources-published-datasets/doi-10-1126science-aaf2403/> and count matrix data and H&E stained brightfield image related to MOB replicate 11 was downloaded.

Code Availability

Instructions to access data, software, and scripts to perform the analysis presented here is available at <https://github.com/MarioniLab/scHOT2019>. scHOT is available as a Bioconductor R package <https://bioconductor.org/packages/scHOT> with detailed vignette available.

References

1. Pijuan-Sala B, et al. A single-cell molecular map of mouse gastrulation and early organogenesis. *Nature*. 2019; doi: 10.1038/s41586-019-0933-9
2. Cao J, et al. The single-cell transcriptional landscape of mammalian organogenesis. *Nature*. 2019; 566:496–502. [PubMed: 30787437]
3. Briggs JA, et al. The dynamics of gene expression in vertebrate embryogenesis at single-cell resolution. *Science*. 2018; 360:1–16.
4. Wagner DE, et al. Single-cell mapping of gene expression landscapes and lineage in the zebrafish embryo. *Science*. 2018; 360:981–987. [PubMed: 29700229]
5. Polioudakis D, et al. A Single-Cell Transcriptomic Atlas of Human Neocortical Development during Mid-gestation. *Neuron*. 2019; doi: 10.1016/j.neuron.2019.06.011
6. Mohammed H, et al. Single-Cell Landscape of Transcriptional Heterogeneity and Cell Fate Decisions during Mouse Early Gastrulation. *Cell Rep*. 2017; 20:1215–1228. [PubMed: 28768204]
7. Mojtahedi M, et al. Cell Fate Decision as High-Dimensional Critical State Transition. *PLOS Biol*. 2016; 14e2000640 [PubMed: 28027308]

8. Basu S, Kumbier K, Brown JB, Yu B. Iterative random forests to discover predictive and stable high-order interactions. *Proc Natl Acad Sci*. 2018; 115:1943–1948. [PubMed: 29351989]
9. Bageritz J, et al. Gene expression atlas of a developing tissue by single cell expression correlation analysis. *Nat Methods*. 2019; 16:750–756. [PubMed: 31363221]
10. Ståh PL, et al. Visualization and analysis of gene expression in tissue sections by spatial transcriptomics. *Science*. 2016; 353:78–82. [PubMed: 27365449]
11. Lubeck E, Coskun AF, Zhiyentayev T, Ahmad M, Cai L. Single-cell in situ RNA profiling by sequential hybridization. *Nat Methods*. 2014; 11:360–361. [PubMed: 24681720]
12. Chen KH, Boettiger AN, Moffitt JR, Wang S, Zhuang X. Spatially resolved, highly multiplexed RNA profiling in single cells. *Science*. 2015; 348
13. Street K, et al. Slingshot: cell lineage and pseudotime inference for single-cell transcriptomics. *BMC Genomics*. 2018; 19:477. [PubMed: 29914354]
14. Campbell KR, Yau C. A descriptive marker gene approach to single-cell pseudotime inference. *Bioinformatics*. 2018; 35:28–35.
15. Haghverdi L, Büttner, Wolf FA, Büttner F, Theis FJ. Diffusion pseudotime robustly reconstructs lineage branching. *Nat Methods*. 2016; 13:845–848. [PubMed: 27571553]
16. Trapnell C, et al. The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. *Nat Biotechnol*. 2014; 32:381–386. [PubMed: 24658644]
17. Van den Berge K, et al. Trajectory-based differential expression analysis for single-cell sequencing data. *Nat Commun*. 2020; 11:1201. [PubMed: 32139671]
18. Lönnberg T, et al. Single-cell RNA-seq and computational analysis using temporal mixture modeling resolves TH1/TH2 fate bifurcation in malaria. *Sci Immunol*. 2017; 2:eaa12192 [PubMed: 28345074]
19. Qiu X, et al. Reversed graph embedding resolves complex single-cell trajectories. *Nat Methods*. 2017; 14:979–982. [PubMed: 28825705]
20. Fuller TF, et al. Weighted gene coexpression network analysis strategies applied to mouse weight. *Mamm Genome*. 2007; 18:463–72. [PubMed: 17668265]
21. Treutlein B, et al. Dissecting direct reprogramming from fibroblast to neuron using single-cell RNA-seq. *Nature*. 2016; 534:391–5. [PubMed: 27281220]
22. Chen S, Mar JC. Evaluating methods of inferring gene regulatory networks highlights their lack of performance for single cell gene expression data. *BMC Bioinformatics*. 2018; 19:232. [PubMed: 29914350]
23. Ghazanfar S, Strbenac D, Ormerod JT, Yang JYHH, Patrick E. DCARS: Differential correlation across ranked samples. *Bioinformatics*. 2018; 35:1–7.
24. Svensson V, Teichmann SA, Stegle O. SpatialDE: Identification of spatially variable genes. *Nat Methods*. 2018; 15:343–346. [PubMed: 29553579]
25. Nitzan M, Karaiskos N, Friedman N, Rajewsky N. Gene expression cartography. *Nature*. 2019; 576:132–137. [PubMed: 31748748]
26. Yang L, et al. A single-cell transcriptomic analysis reveals precise pathways and regulatory mechanisms underlying hepatoblast differentiation. *Hepatology*. 2017; 66:1387–1401. [PubMed: 28681484]
27. Dong J, et al. Single-cell RNA-seq analysis unveils a prevalent epithelial/mesenchymal hybrid state during mouse organogenesis. *Genome Biol*. 2018; 19:31. [PubMed: 29540203]
28. Su X, et al. Single-cell RNA-Seq analysis reveals dynamic trajectories during mouse liver development. *BMC Genomics*. 2017; 18:946. [PubMed: 29202695]
29. Camp JG, et al. Multilineage communication regulates human liver bud development from pluripotency. *Nature*. 2017; 546:533–538. [PubMed: 28614297]
30. Oikawa T, et al. Sall4 Regulates Cell Fate Decision in Fetal Hepatic Stem/Progenitor Cells. *Gastroenterology*. 2009; 136:1000–1011. [PubMed: 19185577]
31. Tanaka M, et al. Mouse hepatoblasts at distinct developmental stages are characterized by expression of EpCAM and DLK1: drastic change of EpCAM expression during liver development. *Mech Dev*. 126:665–76.

32. Sugimoto N, et al. Identification of novel human Cdt1-binding proteins by a proteomics approach: proteolytic regulation by APC/CCdh1. *Mol Biol Cell*. 2008; 19:1007–21. [PubMed: 18162579]
33. Zape JP, Lizama CO, Cautivo KM, Zovein AC. Cell cycle dynamics and complement expression distinguishes mature haematopoietic subsets arising from hemogenic endothelium. *Cell Cycle*. 2017; 16:1835–1847. [PubMed: 28820341]
34. Thakurela S, et al. Gene regulation and priming by topoisomerase II α in embryonic stem cells. *Nat Commun*. 2013; 4:2478. [PubMed: 24072229]
35. Rialland M, Sola F, Santocanale C. Essential role of human CDT1 in DNA replication and chromatin licensing. *J Cell Sci*. 2002; 115:1435–40. [PubMed: 11896191]
36. Lin DM, et al. Spatial patterns of gene expression in the olfactory bulb. *Proc Natl Acad Sci U S A*. 2004; 101:12718–12723. [PubMed: 15304640]
37. Fan XL, Zhang JS, Zhang XQ, Yue W, Ma L. Differential regulation of β -arrestin 1 and β -arrestin 2 gene expression in rat brain by morphine. *Neuroscience*. 2003; 117:383–389. [PubMed: 12614678]
38. Macias M, et al. Spatiotemporal Characterization of mTOR Kinase Activity Following Kainic Acid Induced Status Epilepticus and Analysis of Rat Brain Response to Chronic Rapamycin Treatment. *PLoS One*. 2013; 8:e64455 [PubMed: 23724051]
39. Wilson PO, et al. The immunolocalization of protein gene product 95 using rabbit polyclonal and mouse monoclonal antibodies. *Br J Exp Pathol*. 1988; 69:91–104. [PubMed: 2964855]
40. Gray NW, et al. Dynamin 3 Is a Component of the Postsynapse, Where it Interacts with mGluR5 and Homer. *Curr Biol*. 2003; 13:510–515. [PubMed: 12646135]
41. Kendall RT, et al. Arrestin-dependent Angiotensin AT 1 Receptor Signaling Regulates Akt and mTor-mediated Protein Synthesis. *J Biol Chem*. 2014; 289:26155–26166. [PubMed: 25081544]
42. Girnita L, et al. β -Arrestin is crucial for ubiquitination and down-regulation of the insulin-like growth factor-1 receptor by acting as adaptor for the MDM2 E3 ligase. *J Biol Chem*. 2005; 280:24412–9. [PubMed: 15878855]
43. Bhatnagar A, et al. The Dynamin-dependent, Arrestin-independent Internalization of 5-Hydroxytryptamine 2A (5-HT 2A) Serotonin Receptors Reveals Differential Sorting of Arrestins and 5-HT 2A Receptors during Endocytosis. *J Biol Chem*. 2001; 276:8269–8277. [PubMed: 11069907]
44. Jacque CM, Collet A, Raoul M, Monge M, Gumpel M. Functional maturation of the oligodendrocytes and myelin basic protein expression in the olfactory bulb of the mouse. *Dev Brain Res*. 1985; 21:277–282.
45. Guillemain A, Duchesne R, Crauste F, Gonin-G iraud, Gandrillon O. Drugs modulating stochastic gene expression affect the erythroid differentiation process. *PLoS One*. 2019; 14:e0225166 [PubMed: 31751364]
46. Moris N, et al. Histone Acetyltransferase KAT2A Stabilizes Pluripotency with Control of Transcriptional Heterogeneity. *Stem Cells*. 2018; 36:1828–1838. [PubMed: 30270482]
47. Richard A, et al. Single-Cell-Based Analysis Highlights a Surge in Cell-to-Cell Molecular Variability Preceding Irreversible Commitment in a Differentiation Process. *PLoS Biol*. 2016; 14:e1002585 [PubMed: 28027290]
48. Semrau S, et al. Dynamics of lineage commitment revealed by single-cell transcriptomics of differentiating embryonic stem cells. *Nat Commun*. 2017; 8:1096. [PubMed: 29061959]
49. Stumpf PS, et al. Stem Cell Differentiation as a Non-Markov Stochastic Process. *Cell Syst*. 2017; 5:268–282.e7 [PubMed: 28957659]
50. Wiesner K, Teles J, Hartnor M, Peterson C. Haematopoietic stem cells: entropic landscapes of differentiation. *Interface Focus*. 2018; 8:20180040 [PubMed: 30443337]
51. Lin Y, et al. scMerge leverages factor analysis, stable expression, and pseudoreplication to merge multiple single-cell RNA-seq datasets. *Proc Natl Acad Sci*. 2019; 116:9775–9784. [PubMed: 31028141]
52. Lin Y, et al. Evaluating stably expressed genes in single cells. *Gigascience*. 2019; 8:229815
53. Brennecke P, et al. Accounting for technical noise in single-cell RNA-seq experiments. *Nat Methods*. 2013; 10:1093–5. [PubMed: 24056876]

54. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Series B Stat Methodol.* 1995; 57:289–300.
55. Langfelder P, Zhang B, Horvath S. Defining clusters from a hierarchical cluster tree: the Dynamic Tree Cut package for R. *Bioinformatics.* 2008; 24:719–720. [PubMed: 18024473]
56. Ashburner M, et al. Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nat Genet.* 2000; 25:25–9. [PubMed: 10802651]
57. The Gene Ontology Consortium. The Gene Ontology Resource: 20 years and still GOing strong. *Nucleic Acids Res.* 2019; 47:D330–D338. [PubMed: 30395331]
58. Székely GJ, Rizzo ML, Bakirov NK. Measuring and testing dependence by correlation of distances. *Ann Stat.* 2007; 35:2769–2794.
59. Ghazanfar S, Yang JYH. Characterizing mutation-expression network relationships in multiple cancers. *Comput Biol Chem.* 2016; 63:73–82. [PubMed: 26935398]
60. Pimentel RS, Niewiadomska-Bugaj M, Wang JC. Association of zero-inflated continuous variables. *Stat Probab Lett.* 2015; 96:61–67.

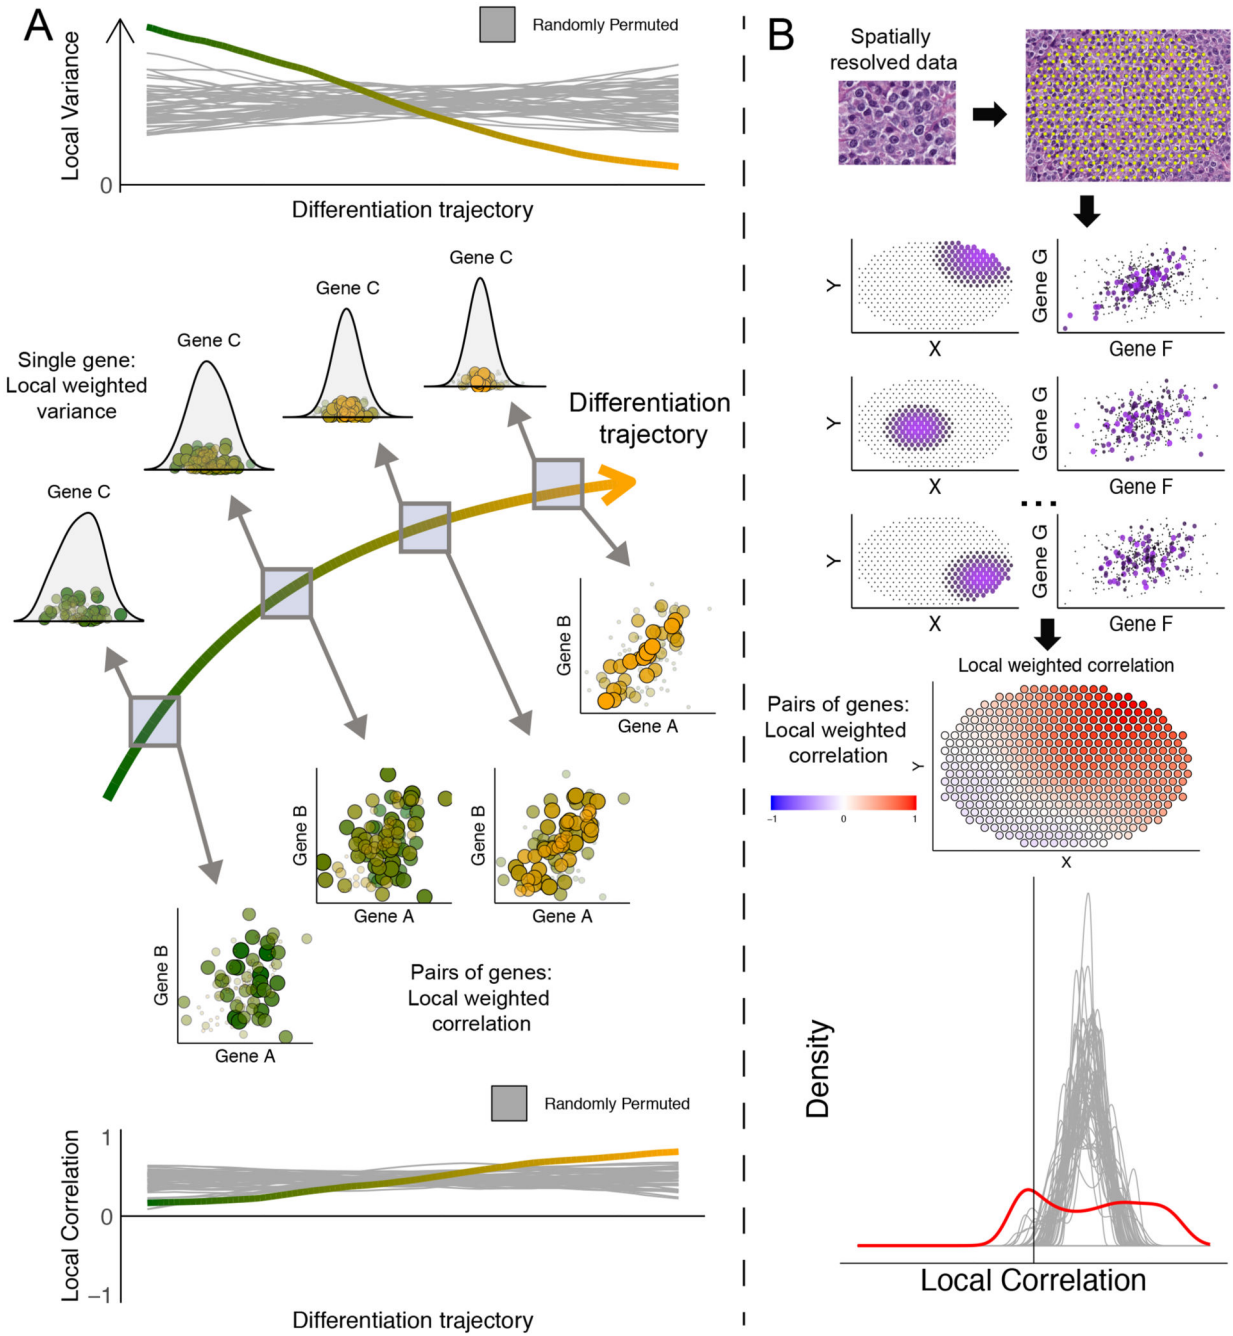


Figure 1. Methods workflow.

A. Example showing a differentiation trajectory where genes are tested for changes in higher order interactions such as variability and correlation along the trajectory. A set of local higher order statistics are calculated, and significance is compared by repeatedly permuting samples (grey curves). The vector of local estimates of higher order statistics are combined using the sample standard deviation to assess how variable it is across time. **B.** Example showing that in a spatial context, sCHOT calculates a field of local estimates of correlation

across space, and compares the variability associated with these with permuted sample points across space.

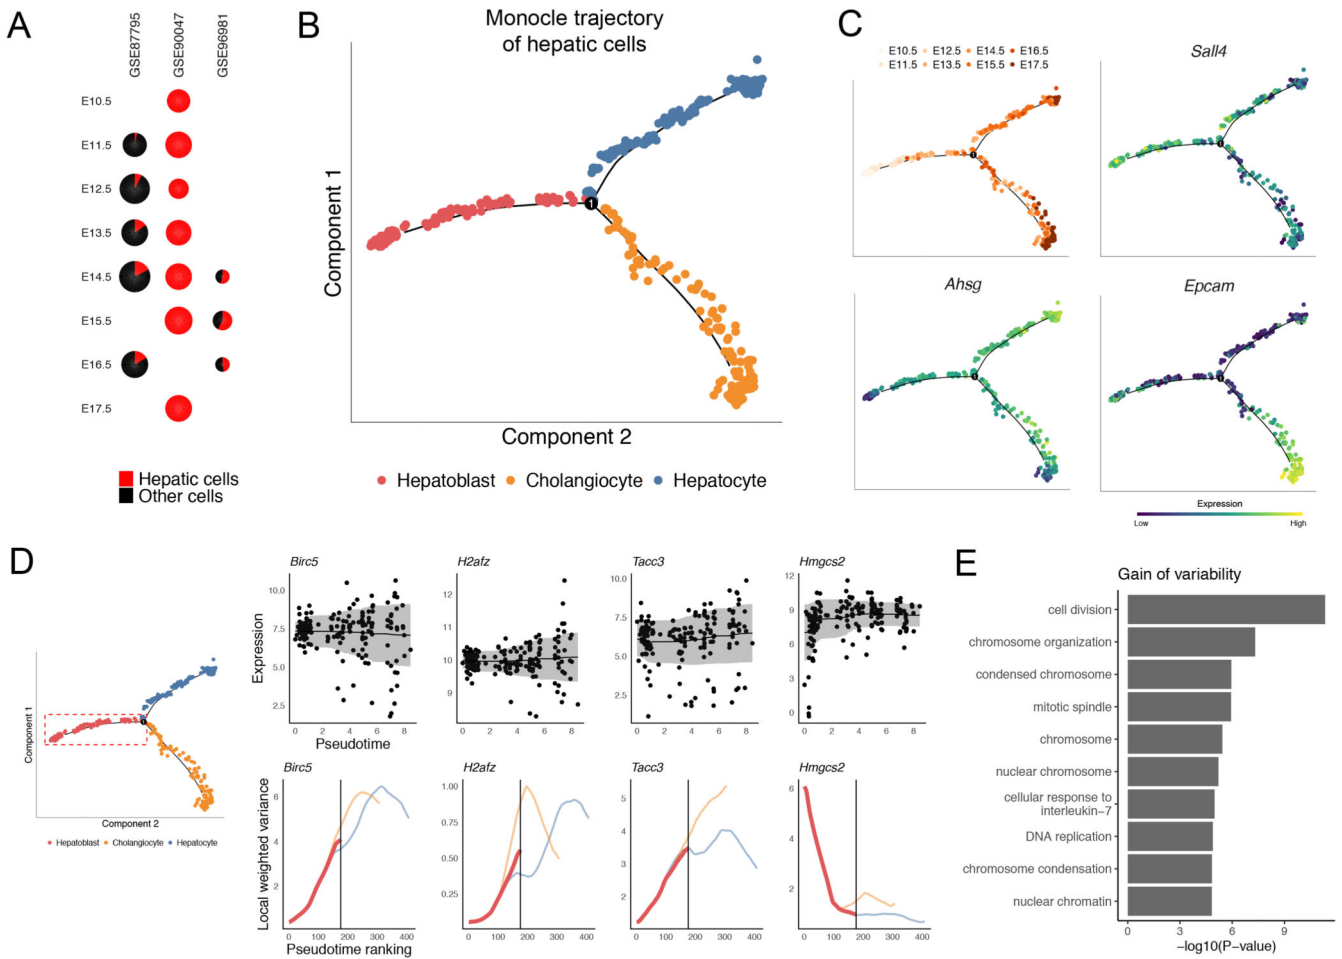


Figure 2. Variability analysis of Developmental Liver Data.

A. Relative proportion of hepatic and non-hepatic cells in the Developmental Liver Data, across original dataset and embryonic stage. **B.** Monocle 2 trajectory of hepatic cells ($n = 540$) showing a bifurcating trajectory of hepatoblasts into either hepatocytes or cholangiocytes. **C.** Panel of embryonic stage for each cell along the differentiation trajectory, and gene expression of markers of each cell type. **D.** A selection of genes significantly associated with a change in variability along the first branch of the differentiation trajectory ($n = 408$) using scHOT with FDR adjusted P-value < 0.1 , scatterplots showing the expression of genes against the pseudotime estimates with shaded ribbons corresponding to adding and subtracting the weighted standard deviation from the weighted mean. Line plots below are of the local estimate of variance for the first branch (thick lines) and further for the two branches (thin lines). Examples are shown of genes that increase in variability via ‘fanning’ of gene expression along the trajectory (*Birc5*), by a skewed distribution arising (*H2afz*), and by changes in the modality of the expression (*Tacc3* from unimodal to bimodal). *Hmgs2* is an example of a gene that decreases in variability. **E.** Gene ontology functional enrichment analysis using one-sided Fisher’s Exact Test of 58 genes that significantly increase in variability along the first branch, grey bar color corresponds to FDR adjusted P-value < 0.05 .

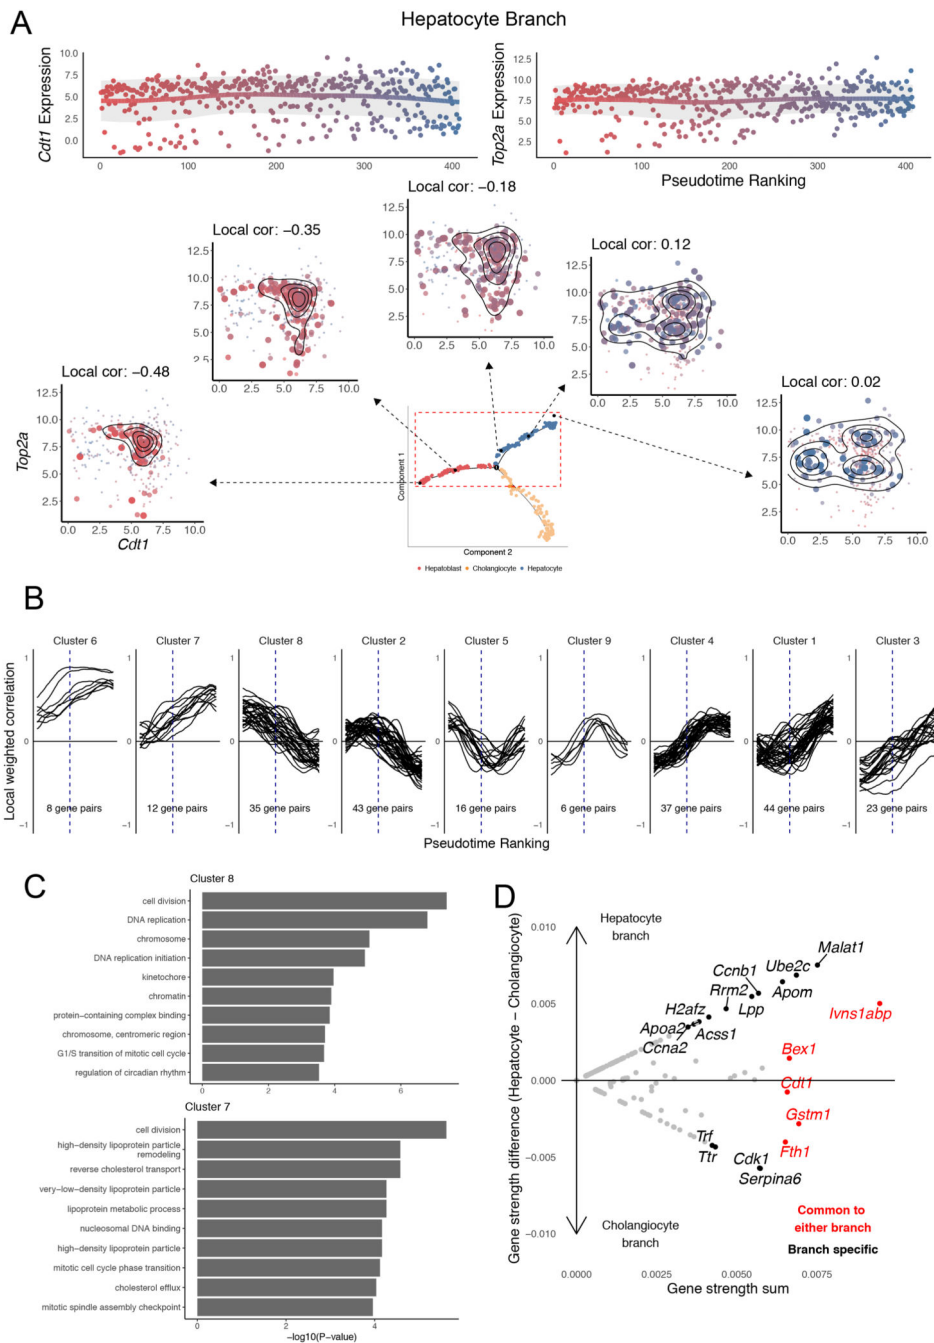


Figure 3. Correlation analysis of Hepatocyte branch of Developmental Liver Data.

A. Sequence of scatterplots showing expression of *Cdt1* and *Top2a* at equally spaced points along the entire trajectory from hepatoblast to hepatocyte ($n = 408$). Points are colored by their position along the trajectory, and point size corresponds to the weight given to that region of the trajectory. Neither gene is significantly differentially expressed or differentially variable along the trajectory using scHOT, but the gene-pair is significantly differentially correlated using scHOT. **B.** Clustering of local weighted correlation using scHOT of all FDR adjusted P -value < 0.2 significant gene-pairs for the hepatocyte branch ($n = 408$), showing

groups of gene-pairs that appear to gain or lose correlation across the trajectory. Vertical dashed line indicates trajectory branchpoint. **C.** Gene ontology functional enrichment analysis using one-sided Fisher's Exact Test of genes belonging to the set of gene-pairs among Clusters 8 and 7 respectively ($n = 47$ genes and $n = 15$ genes respectively) for the hepatocyte branch, grey bar color corresponds to FDR adjusted P-value < 0.05 . **D.** Comparison of hepatocyte and cholangiocyte branches using network strength ($n = 136$ nodes and $n = 71$ nodes respectively) across all significant gene-pairs for either branch. Black labelled genes are significantly branch specific using permutation testing, while red labelled genes are significantly common across both branches (FDR adjusted P-value < 0.05) using permutation testing.

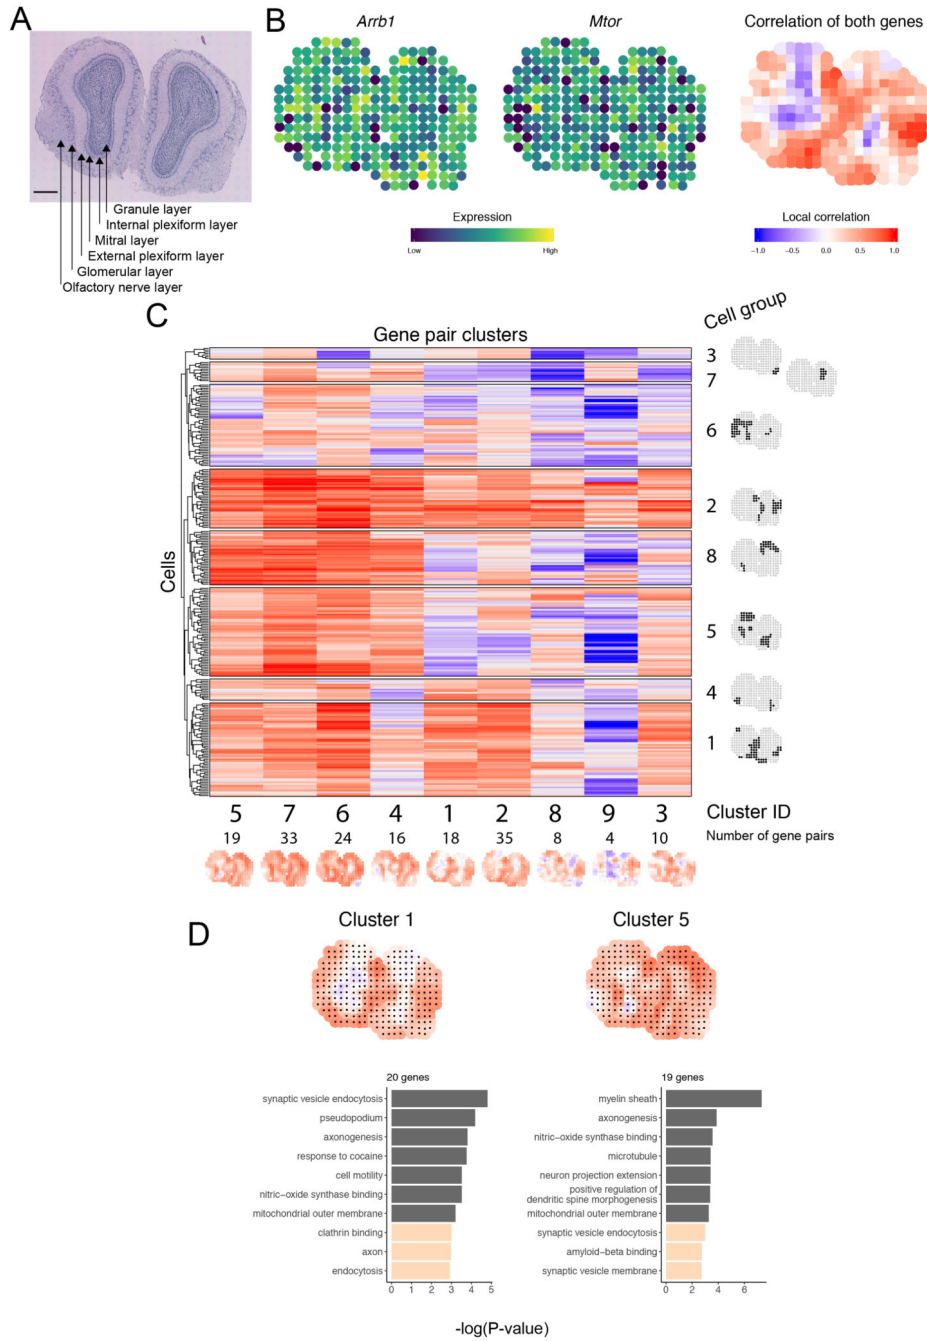


Figure 4. Correlation analysis of Mouse Olfactory Bulb data.

A. An H&E image of mouse olfactory tissue section with labeling of known anatomical layers from Stahl et al¹⁰. Scale bar, 500µm. **B.** Spatial expression plots (n = 262 spatially resolved positions) of two genes, *Arrb1* and *Mtor*, which are not significantly differentially expressed across space using scHOT, but are significantly differentially correlated across space using scHOT with FDR-adjusted P-value < 0.2. The third plot shows the local spatial correlation estimated for these two genes, recapitulating the layered pattern of the olfactory bulb. **C.** Heatmap showing all sampled points (rows) and clusters of significantly spatially

differentially correlated gene-pairs (columns) using schOT, with spatial maps of mean local correlation (bottom row) for each group, and highlighted positions (column on right) for the sampled points grouped into clusters. **D.** Spatial maps of mean local correlation for clusters 1 and 5 (n = 20 and n = 19 genes respectively) of gene-pairs and barplots showing Gene Ontology functional enrichment analysis using one-sided Fisher's Exact Test, grey bar color corresponds to FDR adjusted P-value < 0.05.